arXiv:1612.04335v1 [cs.CV] 13 Dec 2016

Saliency in VR: How do people explore virtual environments?

Vincent Sitzmann[1], Ana Serrano[2], Amy Pavel[3], Maneesh Agrawala[1], Diego Gutierrez[2], Gordon Wetzstein[1]
[1]Stanford University [2]Universidad de Zaragoza [3] University of California Berkeley

Abstract

Understanding how humans explore virtual environments is crucial for many applications, such as developing compression algorithms or designing effective cinematic virtual reality (VR) content, as well as to develop predictive computational models. We have recorded 780 head and gaze trajectories from 86 users exploring omni-directional stereo panoramas using VR head-mounted displays. By analyzing the interplay between visual stimuli, head orientation, and gaze direction, we demonstrate patterns and biases of how people explore these panoramas and we present first steps toward predicting time-dependent saliency. To compare how visual attention and saliency in VR are different from conventional viewing conditions, we have also recorded users observing the same scenes in a desktop setup. Based on this data, we show how to adapt existing saliency predictors to VR, so that insights and tools developed for predicting saliency in desktop scenarios may directly transfer to these immersive applications.

1. Introduction

Virtual and augmented reality systems (VR/AR) provide a new medium that has the potential to profoundly impact our society. The experiences offered by emerging VR/AR systems are not only different from radio, television, or theater, but they are also different from how we experience the real world. For instance, the laws of physics can be broken in VR while maintaining a sense of presence that challenges the user to distinguish between virtual and real.

With unprecedented capabilities for creating immersive content, many questions arise. How does one design or edit 3D scenes effectively? How does one drive user attention in VR? How does one efficiently compress and stream cinematic VR content? To address these and other questions from first principles, it is crucial to understand how users explore virtual environments and to define what constitutes saliency in these immersive applications. A detailed understanding of visual attention in VR would not only help answer the above questions, but also inform future designs of user interfaces, eye tracking technology, and other key aspects of VR systems. A crucial requirement for developing this understanding, however, is access to behavioral data.

We have implemented a head-mounted display (HMD) system that allows us to record users' gaze direction as well as their head orientation (yaw, pitch, and roll). The recorded data is matched to the presented visual stimuli for subsequent analysis. We have recorded 780 head and gaze trajectories from 86 people in 22 static virtual environments, represented as omni-directional stereo panoramas, starting from 4 different viewpoints. Data is recorded from users in a standing position using the HMD; we also capture data of users observing the same scenes on a desktop monitor for comparison. Figure 1 shows panoramic views of all scenes with superimposed saliency for the VR condition. Salient image regions in these 360° views match characteristics predicted by previous studies on saliency with conventional displays (e.g., [2, 5, 12]): people fixate consistently on faces, text, and vanishing points. However, view behavior in virtual environments is governed by the interaction of head orientation, gaze, and other kinematic constraints, which make these immersive scenarios more complex than observing a conventional display.

To further our understanding of view behavior and saliency in VR, we capture and analyze gaze and head orientation data for consistency of fixations, correlation between perceptual aspects such as gaze velocity or acceleration and scene entropy, as well as interactions between gaze and head orientation. Using the recorded data, we demonstrate that existing saliency models can be adapted to predict saliency for VR, as well as other factors that give insights into how people explore virtual environments. We will make all captured data public, to foster future research on this topic.

2. Related Work

Modeling human gaze behavior and predicting visual attention has been an active area of vision research. In their seminal work, Koch and Ullman [21] introduced a model for predicting salient regions from a set of image features. Motivated by this work, many models of visual attention have been proposed throughout the last three decades. Most of these models are based on bottom-up, top-down, or hy-

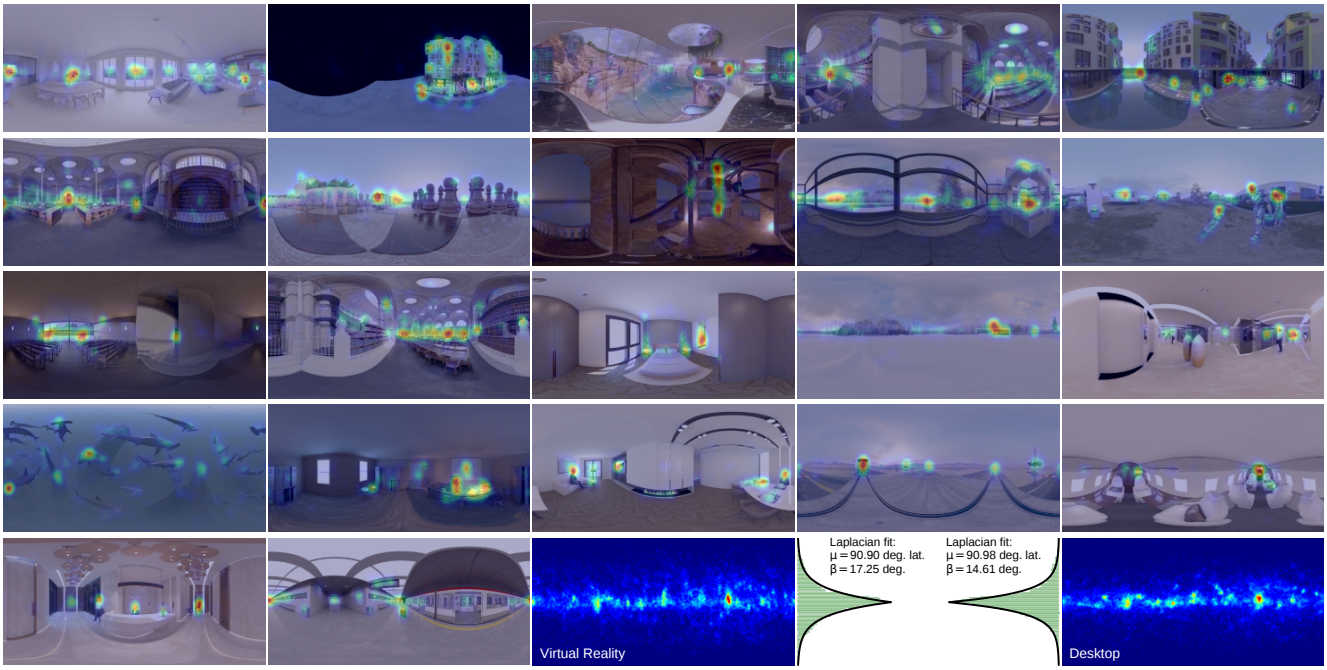

Figure 1: A mosaic of all 22 panoramas used for our study with superimposed saliency levels in an equirectangular projection. All of these saliency maps are computed from data captured in the *VR condition*, but we show average saliency maps for both the VR and the *desktop condition* in the bottom right. These average maps demonstrate an "equator bias" that is well-described by a Laplacian fit modeling the probability of a user fixating on an object at a specific latitude.

brid approaches. Bottom-up approaches build on a combination of low-level image features, including color, contrast, or orientation [17, 29, 7, 20] (see Zhao and Koch [40] for a review). Top-down models take higher-level knowledge of the scene into account such as context or specific tasks [18, 19, 15, 28, 36]. Recently, advances in machine learning and particularly convolutional neuronal networks (CNNs) have fostered the convergence of top-down and bottom-up features for saliency prediction, producing more accurate models [37, 41, 38, 32, 27]. Bylinskii et al. [4] explored the shortcomings of state-of-the-art saliency models and provided a rigorous base for saliency map benchmarking. Recent work also attempts to extend CNN approaches beyond classical 2D images by computing saliency in more complex scenarios such as stereo images [16, 8] or video [6, 26]. Building on the rich literature in this area, we explore user behavior and visual attention in immersive virtual environments.

What makes VR different from desktop viewing conditions is the fact that head orientation is used as a natural interface to control perspective. The interaction of head and eye movements are complex and neurally coupled, for example via the vestibulo-ocular reflex [23]. For more information on user behavior in VR, we refer to Ruhland et al. [33], who provide a review of eye gaze behavior, and Freedman [13], who discusses the mechanisms that characterize the coordination between eyes and head during visual orienting movements. Also, Kollenberg et al. [22] propose a basic analysis of head-gaze interaction in VR. Closer to our research is the recent work of Nakashima et al. [30], who propose to use a head direction prior to improve accuracy in saliency-based gaze prediction through simple multiplication of the gaze saliency map by a Gaussian head direction bias. The data collected in this paper and in-depth analyses augment prior work in this field and may allow for future data-driven models for visual behavior to be learned.

Finally, gaze tracking has found many applications in VR user interfaces [35] and gaze-contingent displays [11], such as foveated rendering. The understanding of user behavior we aim to develop with our work could influence future work in all of those areas. If successful, gaze prediction in VR could make costly eye trackers expendable.

3. Dataset of human behavior in VR

3.1. Data capture

Conditions We recorded standing users who observed 22 omni-directional stereo panoramas (see Fig. 1), under two different conditions: using head mounted displays, and viewing the same scenes on a desktop monitor. In the desktop condition, the scenes are monoscopic, and users navigate with a mouse, instead of using head rotation to explore

the scenes. All scenes are computer-generated by artists; we received permission to use them for this study. For each scene, we recorded four different starting points, spaced at 90° longitude, which results in a total of 88 test conditions. These starting points were chosen to cover the entire longitudinal range, while keeping the number of different conditions tractable.

Participants For experiments in VR, 86 users participated in our study (68 male, 18 female, age 17-22). Users were asked to first perform a stereo vision (Randot) test to quantify their stereo acuity. For desktop experiments, we recruited 44 additional participants (27 male, 17 female, age 18-33). For both tests, all participants reported normal or corrected-to-normal vision.

Procedure All VR scenes were displayed using an Oculus DK2 head-mounted display, equipped with a pupil-labs[1] stereoscopic eye tracker recording at 120 Hz. The DK2 offers a field of view of $106 \times 87°$. Prior to displaying test scenes, the eye tracker was calibrated using a per-user, per-use procedure. The Unity game engine was used to display all scenes and record head orientation while the eye tracker collected gaze data on a separate computer. Users were instructed to freely explore the scene and were provided with a pair of earmuffs to avoid auditive interference in the visual task. To guarantee the same starting condition for all users, a gray environment with a small red box was displayed between different scenes, and users were instructed to find it. 500 ms after they aligned their head direction with the red box, a new scene would appear. Each panorama was displayed for 30 seconds. Scenes and starting points were randomized such that each user would only see the same scene once from a single random starting point. Each user was shown 8 scenes; the total time per user experiment, including calibration, was approximately 10 minutes.

For the desktop condition, users sat 0.45 meters away from a 17.3" monitor with a resolution of 1920×1080 px, covering a field of view of $23 \times 13°$. We used a Tobii EyeX eye tracker with an accuracy of 0.6° at a sampling frequency of 55 Hz [14]. The image viewer displayed a rectilinear projection of a $97 \times 65°$ viewport of the panorama, typical for desktop panorama viewers on the web. We calibrated the eye tracker and instructed the users how to use the image viewer, before showing the 22 scenes for 30 seconds each. We paused the study after half of the scenes to recalibrate the eye tracker. In this condition, we only collected gaze data but not head orientation because the field of view is much smaller and users rarely re-orient their head. Instead, we recorded where the users interactively place the virtual camera in the panorama as a proxy for head orientation.

[1] https://pupil-labs.com

3.2. Data processing

Processing eye tracker samples For the VR setup, we linearly interpolated all measurements when the eye tracker reported a confidence below 0.9. Since the eye tracker and the head position tracker had different sampling rates, each head measurement was assigned to the gaze measurement with the closest timestamp. For the desktop condition, the eye tracker yielded a single gaze position. We linearly interpolated low-confidence measurements and matched the higher-frequency eye tracking data with the lower-frequency camera orientation data of the desktop condition. For both conditions, we calculated head/camera and gaze speeds, as well as accelerations as the first and second derivatives of the latitude and longitude using the forward finite difference method. Before computing head and gaze statistics, we manually identified and filtered outliers at least 5 standard deviations away from the mean, accounting for less than 0.5% of measurements.

Fixations To identify fixations, we transformed the normalized gaze tracker coordinates to latitude and longitude in the 360° panorama. This is necessary to detect users fixating on panorama features while turning their head. We used thresholding based on dispersion and duration of the fixations [34]. For the VR experiments, we set the minimum duration to 150 ms [34] and the maximum dispersion to 1° [1]. For the desktop condition, the Tobii EyeX eye tracker output showed significant jitter, and visual inspection revealed that this led to skipped fixations. We thus smoothed this data with a running average of 2 samples and detected fixations with a dispersion of 2°. We counted the number of fixations at each pixel location in the panorama. Similar to Judd et al. [19], we only consider measurements from the moment where user's gaze left the initial starting point to avoid adding trivial information. We convolved these fixation maps with a Gaussian with a standard deviation of 1° of visual angle to yield continuous saliency maps [24].

4. Analysis

Comparison metrics Following previous works, we use three main metrics to compare saliency maps. We use the Inter-Observer Visual Congruency (IOVC) metric [25] when investigating whether two groups of users fixated in the same regions of a scene. This is achieved by generating a saliency map with one set of fixations, and measuring the percentage of the other set of fixations that fall in the top 25% most salient regions of that map and vice versa, then averaging the two percentages. An IOVC of 100% thus indicates a strong overlap of salient regions, while 0% indicates no overlap. The IOVC metric is robust to measurement error in the exact position of the fixations. For comparisons between continuous saliency maps, we use the Pear-

son correlation (CC) metric, and the Earth's Mover Distance (EMD), metric following the work of Bylinskii et al. [4]. A high CC score is achieved when two saliency maps present a strong overlap, and the magnitudes of the saliency values are similar. We show all the results with the CC metric throughout this section while similar results with the EMD metric can be found in the supplementary material.

Equator bias In both VR and desktop conditions, we have identified an effect that is similar to the center-bias described by Nuthmann and Henderson [31] for images displayed on a traditional screen. Users tend to fixate around the equator of the panoramas, with very few fixations in latitudes far from the equator. To quantify this *equator bias*, we first calculate the average of all 22 saliency maps, and filter out the fixations that users made until they left the close vicinity (20° longitude) of the starting point. We then marginalize out the longitudinal component of the saliency map, and fit a Laplace distribution with location parameter μ and diversity β, to the latitudinal component. This particular distribution empirically matched the data best among several tested distributions. Figure 1 depicts the average saliency map, as well as the Laplace fit to the latitudinal distribution and its parameters, for both the VR and the desktop experiments. While the mean is the same for the two viewing conditions, the desktop condition equator bias has a lower diversity. This Laplacian can be used to generate a baseline for saliency maps in VR, which we refer to as the "equator bias baseline". This baseline achieves a mean CC of 0.33 with the collected saliency maps, and will be used in Section 5 when evaluating other saliency predictors.

We note that all scenes, except for one, have a clear horizon line, around which most relevant objects of the scenes are located. Therefore, the observed bias could be a result of this specific dataset. Nevertheless, the horizon line is a characteristic shared with most virtual environments, and also the real world, which makes saliency in VR different from natural images used by other saliency studies.

Consistency of fixations We found that for some scenes, users fixate in the same locations, while in others users' fixations are scattered all over the scene. Here, we show results for the VR viewing condition; please find the same analysis for the desktop condition in the supplement. Following a similar procedure to Judd et al. [19] we analyze the *consistency of human fixations* by computing the Shannon entropy of the saliency map as $-\sum_{i=1}^{N} s_i^2 log(s_i^2)$, with s being the saliency map and N the number of pixels, and then normalizing it. Figure 2 shows the saliency maps of the scenes with highest and lowest entropy.

Additionally, we analyze the performance of human fixations as predictors of saliency or *human performance*. We use a *receiver operating characteristic* curve (ROC) metric [4] to calculate the ability of the i^{th} user to predict the ground truth saliency map computed from the fixations of all the other users averaged. A single point in the ROC curve is computed by finding the top $n\%$ most salient regions of the ground-truth saliency map (leaving out the i^{th} user), and then calculating the percentage of fixations of the i^{th} user that fall into these regions. We show the ROC for all the 22 scenes in Figure 2 second column (middle).

We further analyze the correlation between human performance and entropy. We test the entropy against the accuracy of the human performance on the 10% most salient regions of the scene (dotted line in the figure) since in our scenes the regions of interest lie mostly within those 10%. We calculate Pearson and Kendall correlation coefficients (Figure 2 right column, bottom). Pearson correlation tests the linear relationship between both variables, while

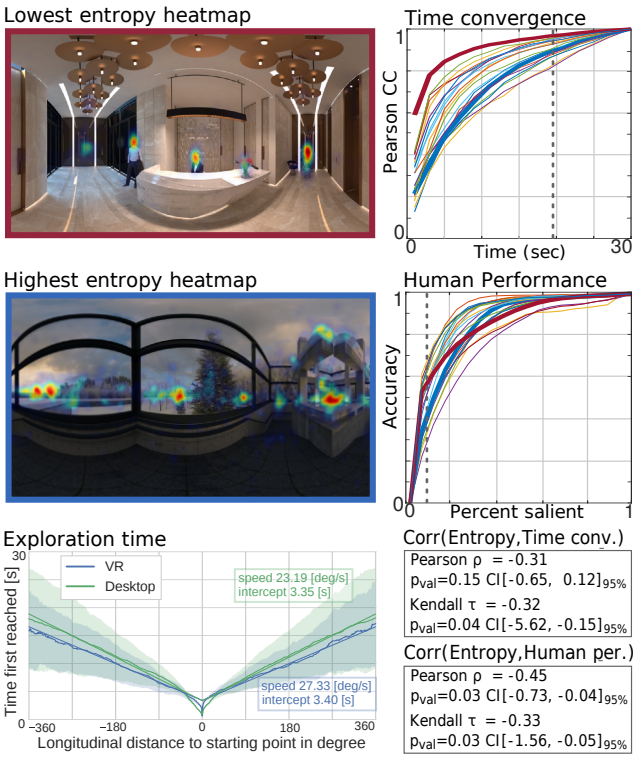

Figure 2: The first column shows the saliency maps with highest (top) and lowest (middle) entropy in our dataset. Exploration time (bottom left) is the average time until a specific longitudinal offset from the starting point is reached. The second column shows the temporal convergence of the saliency map for each scene (top) computed as the similarity between the saliency maps at each time step and the fully-converged saliency maps; the ROC curve of human performance averaged across users (middle); and the correlations (bottom) of both with the entropy (Pearson linear correlation and Kendall rank correlation).

Kendall performs a rank correlation. To measure statistical significance we calculate p-values, as well as 95% confidence intervals (CI). Results show a statistically significant inverse correlation for both Pearson and Kendall coefficients. These results suggest that fixations in scenes with higher entropy are not only very scattered, but also that fixations are very inconsistent across users.

Exploration time and temporal convergence We evaluate the speed with which users explore a given scene. In Figure 2 (bottom left) we show the *exploration time*, which is the average time that users took to move their eyes to a certain longitude relative to their starting point, for both the VR and desktop conditions. We shade the zone of one standard deviation below and above both curves. Exploration times for the two viewing conditions are very similar. On average, users can be expected to having fully explored the scene after about 19 seconds.

We further evaluate the temporal convergence of the saliency maps, i.e. after how much time of exploration the resulting fixation map has converged to the ground-truth saliency map. For every scene, we calculate the saliency map at different time steps, and compute the similarity between these and the fully-converged saliency map. We show the temporal evolution of the CC score for all the scenes in Figure 2 second column (top). Again, we compute Pearson and Kendall correlation coefficients, this time testing the entropy against the convergence after 19 seconds (dotted line in the figure). According to the *exploration time* we assume that after this time users will have already traveled through the full scene at least once. Results indicate no statistical significance for the linear test, however there is a statistically significant inverse rank correlation between the two factors. This leads to the conclusion that for scenes with high entropy the saliency maps take longer to converge. These results are aligned with the intuitive notion that users explore scenes with few and clear salient regions faster, fixating very quickly in such regions; for more complex scenes, users wander around the scene and fixations take longer to converge (in addition to being very scattered).

Comparing VR and desktop saliency maps We compare saliency maps obtained from the VR and desktop experiments. Visual inspection shows a high similarity between the saliency maps (see supplement). This is confirmed by a high mean IOVC score of 80.68, which is computed by averaging over all per-scene pairwise IOVC scores between saliency maps of the two conditions. This is encouraging, as it allows using a desktop environment to collect saliency maps very close to the VR saliency maps. Since desktop experiments are much easier to control, this insight makes it more feasible to collect training sets for data-driven saliency prediction in future VR systems.

Figure 3: Ground truth saliency (left), and cropped saliency maps for four different starting points (right). IOCV values are mean of pairwise comparisons of one map to the three others. Visual inspection and IOCV scores suggest little influence of the starting point on the final saliency map.

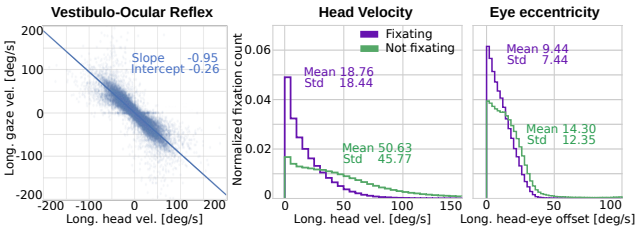

Figure 4: Left: the vestibulo-ocular reflex demonstrated by an inverse linear relationship of gaze and head velocities. Middle and right: distributions of longitudinal head velocity and longitudinal eye eccentricity while fixating or not.

Influence of the starting point We evaluate the influence of the initial starting point in the converged saliency map for both the VR and desktop conditions. In Figure 3 we show the ground truth saliency map for a scene (left), and cropped regions of the saliency map generated when starting at four different, equally-spaced starting points. For both conditions, we generate saliency maps for each scene and starting point. Visual inspection reveals that users starting from different points investigate the same regions within the experiment (see supplement). We find that regions that are salient in the complete saliency map are always also investigated in each of the starting point saliency maps, though they sometimes receive different degrees of attention. We compute the IOVC metric for each scene and each pair of viewpoints. On average, 99.84% of fixations starting from one point lie within the top 25% most salient regions in the saliency map from the other starting point, and vice versa. Desktop results are similar with 93.18% IOVC. We can thus conclude that the starting condition has little impact on which regions of the scene are attended to over the course of 30 seconds.

Head and gaze statistics For the VR viewing condition, the mean duration of fixations across scenes is 260 ms $\pm$ 126. For the desktop condition, we measure 245 ms $\pm$ 114. Both are in the range reported for traditional screen view-

	Mean [°/s]	99_{th} per. [°/s]
Head speed (lat, lon)	(7.3, 24.9)	(74.2, 151.6)
Gaze speed (lat, lon)	(31.1, 44.6)	(295.2, 359.2)
Camera speed (lat, lon)	(20.5, 32.1)	(325.8, 562.5)
Gaze speed (lat, lon)	(38.2, 64.3)	(252.2, 467.1)

Table 1: Mean and standard deviation of head and gaze velocities for the VR condition (top); camera and gaze velocities for the desktop condition (bottom).

ing conditions [34]. In VR, the mean eye rotation relative to the head across scenes is 12.17° ± 10.76 which is consistent with the analysis performed by Kollenberg et al. [22], reporting narrower eye-rotations for HMD than for traditional viewing conditions. The mean number of fixations is 49.83 ± 13.04 for VR, and 40.04 ± 13.36 for the desktop condition. In table Table 1, we show the mean speed for gaze and head movements (longitudinal and latitudinal) for VR, as well as the mean gaze speed for the desktop condition. Interestingly, both head and gaze move much slower in the latitudinal direction.

We further analyze interactions between eye and head rotation (latitude and longitude), velocity, and acceleration. To analyze the interaction between eye and head when shifting to a new target, we offset head and gaze acceleration measurements relative to each other in time and compute the cross-correlation for different temporal shifts. Our data reveals that head follows gaze with an average delay of 58 ms, where the largest cross-correlation is observed. This is consistent with previous works [13, 10] reporting delays in head movement when shifting to a non-predictable target (i.e., not trained nor premeditated gaze shifts). We have also identified the *vestibulo-ocular reflex* [23] in our dataset, where gaze compensates for head movements while fixating: Figure 4 (left) shows the expected inverse linear relationship between head velocity and relative gaze velocity when fixating. We also found that the statistics of head and eye movements differ when users fixate versus when they do not fixate. Figure 4 (middle) demonstrates that users move their head at velocities in longitudinal direction significantly below the average head speed when they are fixating. As expected, when users are not fixating, head speeds are above average. The same effect, but to a lesser degree, can be seen in the latitudinal head velocity (see supplement). Further, Figure 4 (right) shows that the longitudinal rotation angle of the eyes relative to the head orientation (eye eccentricity) is significantly smaller when users are not fixating. The same, however, does not seem to be true for the eye offset in latitudinal direction (see supplement). An interesting conclusion is that users appear to behave in two different modes: *exploration* and *re-orientation*. Eye fixations happen in the exploration mode, when users have "locked in" on a salient part of the scene, while movements to new, salient regions happen in the re-orientation mode.

Head orientation as gaze predictor We analyze whether longitudinal head velocity at every time step is an indicator of fixations, and in particular if saliency maps obtained by thresholding longitudinal head velocity can sensibly be used as approximations of gaze saliency maps. Following Figure 4 (middle), we assume that whenever head speed falls below the threshold of 19 $^\circ/s$, the likelihood that the user is currently fixating is high. We create a head orientation saliency map by counting the number of measurements that fulfill this criterion at each pixel location. Following Figure 4 (left), we then blur the obtained saliency map with a Gaussian kernel of size 9.0° of visual angle. This accounts for the mean eye offset. Visual inspection (see supplement) shows that this approximation captures the approximate location of highly salient regions. The computed saliency maps achieve an average CC of 0.51, compared to a CC of 0.33 for the equator bias baseline. We conclude that head saliency maps, which are easily obtainable with inertial measurement units, can be a valuable tool to analyze the approximate regions that users attend to in a scene without the need for additional eye-tracking hardware.

5. Predicting Saliency in VR

In this section, we show how existing saliency prediction models can be adapted to VR. We further ask whether the problem of time-dependent saliency prediction is a well-defined one that can be answered with sufficient confidence.

5.1. Predicting Converged Saliency Maps

Saliency prediction is a well-explored topic, with many existing models evaluated by the MIT Saliency Benchmark [3]. Given an image, these models predict a saliency map in the form of a 2D probability distribution which characterizes the probability that a location is fixated within a two-second timeframe. The MIT benchmark and most existing prediction models, however, assume that users sit in front of a screen while observing the images – ground truth data is collected by eye trackers recording precisely this behavior. We argue that observing 360° panoramic image content in VR may be different enough to warrant closer investigation. VR is different from conventional screens in that users naturally use both head orientation and gaze to visually explore scenes. In addition, kinematic constraints, such as sitting in a non-swivel chair, may affect user behavior.

The datasets we recorded may be insufficient to train new data-driven behavioral models, but they may be sufficient to study the differences between viewing behavior in VR and with a conventional screen. Ideally, we would like to leverage existing saliency prediction models for VR scenarios. In this context, two primary challenges arise: (i) mapping a 360° panorama to a 2D image (the required input for existing models) distorts the content due to the projective mapping from sphere to plane; and (ii) head-gaze

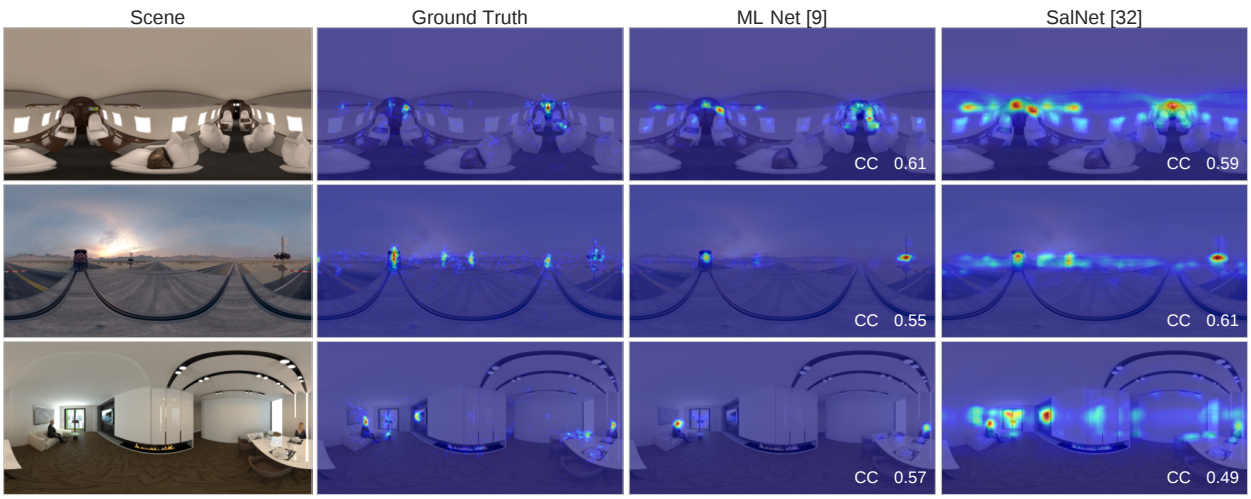

Figure 5: Time-independent saliency prediction for omni-directional stereo panoramas can leverage existing saliency models. For this purpose, the target panorama (left) is divided into small patches (we use $60\times65°$ per patch). Each of these patches can be represented with minimal perspective distortions and is processed with a saliency model of choice. The resulting patches are stitched together into a panorama and weighted by the latitudinal Laplacian function representing the equator bias (center right and right columns). Both quantitatively and qualitatively, in many cases this simple procedure achieves reasonable results compared to the ground truth saliency maps recorded with our gaze tracker in VR (second column).

interaction may require special attention for saliency prediction in VR. We address both of these issues in an intuitive way. First, we split the target panorama into several patches, each exhibiting minimal projective distortion so they can be processed separately by saliency predictors. Second, we weight the resulting saliency panorama by the equator bias derived in Section 4 to account for longitudinal viewing preference. This simple approach is independent of any particular saliency predictor, and implemented as follows:

1: extract gnomonic projection patches from a panorama
2: apply a 2D saliency predictor to each patch
3: stitch patches back into a saliency panorama
4: weight saliency panorama by equator bias
5: normalize resulting saliency panorama

We evaluate this procedure both quantitatively and qualitatively. Table 2 lists mean and standard deviation of Pearson CC for all 22 scenes in the VR scenario, and for users exploring the same scenes in the desktop condition. These numbers allow us to analyze how good and how consistent across scenes a particular predictor is. We test the equator bias by itself as a baseline as well as the two highest-ranked models in the MIT benchmark where source code is available: ML-Net [9] and SalNet [32]. We see that the two advanced models perform very similar, but both much better than the equator bias alone. We also see that both of these models predict viewing behavior in the desktop condition better than is the case for the VR condition. This is intuitive, because the desktop condition is the condition the models were trained for originally.

	Horizon Bias	ML-Net [9]	SalNet [32]
VR standing	$\mu=0.33, \sigma=0.13$	$\mu=0.45, \sigma=0.11$	$\mu=\mathbf{0.45}, \sigma=0.13$
Desktop	$\mu=0.43, \sigma=0.12$	$\mu=\mathbf{0.58}, \sigma=0.11$	$\mu=0.57, \sigma=0.12$

Table 2: Quantitative comparison of predicted saliency maps using a simple equator bias, and two state-of-the-art models. Numbers show average mean and standard deviation of Pearson correlations, for each scene, between prediction and ground truth recorded from users exploring 22 scenes in the VR and desktop conditions.

In Figure 5 we compare the saliency maps of three scenes recorded under the VR condition (all scenes in the supplement). Qualitatively, both predictors perform reasonably well. Nevertheless, we believe that viewing behavior in VR may be driven by higher-level tasks and cognitive processes, possibly more so than observing conventional images. Current-generation saliency predictors do not take this into consideration. We hope that our dataset, along with our simple evaluation metric outlined above, will be helpful for benchmarking future VR saliency predictors.

5.2. Can Time-dependent Saliency Maps be Predicted with Sufficient Confidence?

VR scenes dictate viewing conditions much different from those assumed in models for the classic visual saliency problem. Apart from the perspective distortion problem discussed above, the question of temporal development arises: for users starting to explore the scene at a given starting

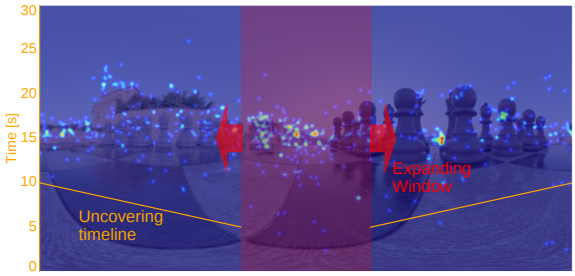

Figure 6: Time-dependent saliency prediction by uncovering the converged saliency map with the average exploration speed determined in Section 4.

point, how can we describe the probability that they fixate at specific coordinates at a time t? We use the insights from Section 4 to build a simple baseline model for this problem. Figure 2 (bottom left) yields an estimate for when users reach a certain longitude on average. We thus model the time-dependent saliency map of a scene with an initially small window that grows larger over time to progressively uncover more of a converged (predicted or ground-truth) saliency map. The part of the saliency map within this window is the currently active part, while the parts outside this window are set to zero. The left and right boundaries of the window are widened with the speed predicted in Figure 2.

Figure 6 visualizes this approach. We generate the time-dependent saliency maps for all 22 scenes and compare them with ground truth. We use the fully-converged saliency map as a baseline. The predicted, time-dependent saliency maps model the recorded data better than the converged saliency map within the first 6 seconds. Subsequently, they perform slightly worse until the converged map is fully uncovered after about 10 seconds, and the model is thus identical to the baseline. Our simple time-dependent model achieves an average CC over all scenes, viewpoints and the first 10 seconds of 0.57 (uncovering the ground-truth saliency map), while using the converged saliency map as a predictor yields a CC of just 0.47.

Although this is useful as a first-order approximation for time-dependent saliency, there is still work ahead to adequately model time-dependent saliency over prolonged periods. In fact, due to the high inter-user variance of recorded scan paths, the problem of predicting time-dependent saliency maps may not be a well-defined one. Perhaps a real-time approach that would use head orientation measured by an inertial measurement unit to predict where a specific user will look next could be more useful than trying to predict time-dependent saliency without any knowledge of the specific user. We leave this investigation for future work.

6. Discussion

In summary, we collect a dataset that records gaze and head orientation for users observing omni-directional stereo panoramas in VR. We also capture users observing the same scenes in a desktop scenario, exploring the panoramas with mouse-based interaction. We will make all this data available to the public.

The primary insights of our study are: (1) saliency in VR seems to be in good agreement with saliency in conventional displays; as a consequence, existing saliency predictors can be applied to VR using a few simple modifications described in this paper. (2) head and gaze interaction are coupled in natural or VR viewing conditions, but they differ when users fixate and when they do not – the collected data may enable models to be learned that could approximate gaze movement from head and image information alone, without costly eye trackers. (3) while our simple time-dependent model can approximate where people look within the first few seconds after entering a new scene, accurate prediction is not possible for more than a few seconds at the moment.

These insights could have a direct impact on a range of common tasks in VR. Adaptive compression and streaming of omni-directional stereo content [39], for example, would benefit from an understanding of what people are likely to look at in VR. Placing cuts in dynamic cinematic VR content adaptively in runtime could ensure that users see what the artists would like them to see. Placing information or advertisement into virtual environments could be more effective with a deeper understanding of saliency in VR. Finally, VR allows for unprecedented types of behavioral data to be collected and analyzed, which would be useful for basic cognitive science experiments and for learning behavior.

Future Work The collected dataset of 780 head and gaze trajectories for 22 omni-directional stereo panoramas may allow predictive models for head and gaze trajectories to be learned. It would be interesting to explore how such models could improve low-cost but imprecise gaze sensors, such as electrooculograms. Finally, future work could extend the data collection and analysis to videos and multimodal experiences that include audio.

7. Acknowledgements

The authors would like to thank Belen Masia for fruitful discussions and insights and Jaime Ruiz-Borau for support with experiments. This research has been partially funded by an ERC Consolidator Grant (project CHAMELEON), the Spanish Ministry of Economy and Competitiveness (projects LIGHTSLICE and LIGHTSPEED), the NSF/Intel Partnership on Visual and Experiential Computing (NSF IIS 1539120) as well as the Intel Compressive Sensing Al-

liance. Ana Serrano was supported by an FPI grant from the Spanish Ministry of Economy and Competitiveness. Gordon Wetzstein was supported by a Terman Faculty Fellowship and an Okawa Research Grant.

We thank the following artists, photographers, and studios who generously contributed their omni-directional stereo panoramas for this study: Dabarti CGI Studio, Attu Studio, Estudio Eter, White Crow Studios, Steelblue, Blackhaus Studio, immortal-arts, Chaos Group, Felix Dodd, Kevin Margo, Aldo Garcia, Bertrand Benoit, Jason Buchheim, Prof. Robert Kooima, Tom Isaksen (Charakter Ink.), Victor Abramovskiy (RSTR.tv)

References

[1] P. Blignaut. Fixation identification: The optimum threshold for a dispersion algorithm. *Attention, Perception, & Psychophysics*, 71(4):881–895, 2009. 3

[2] A. Borji and L. Itti. State-of-the-art in visual attention modeling. *IEEE Trans. PAMI*, 35(1):185–207, 2013. 1

[3] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba. Mit saliency benchmark. 6

[4] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand. What do different evaluation metrics tell us about saliency models? *arXiv preprint arXiv:1604.03605*, 2016. 2, 4

[5] Z. Bylinskii, A. Recasens, A. Borji, A. Oliva, A. Torralba, and F. Durand. Where should saliency models look next? In *Proc. ECCV*, pages 1–16, 2016. 1

[6] S. Chaabouni, J. Benois-Pineau, O. Hadar, and C. B. Amar. Deep learning for saliency prediction in natural video. *arXiv preprint arXiv:1604.08010*, 2016. 2

[7] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. Torr, and S.-M. Hu. Global contrast based salient region detection. *IEEE Trans. PAMI*, 37(3):569–582, 2015. 2

[8] R. Cong, J. Lei, C. Zhang, Q. Huang, X. Cao, and C. Hou. Saliency detection for stereoscopic images based on depth confidence analysis and multiple cues fusion. *IEEE Signal Processing Letters*, 23(6):819–823, 2016. 2

[9] M. Cornia, L. Baraldi, G. Serra, and R. Cucchiara. A deep multi-level network for saliency prediction. In *Proc. ICPR*, 2016. 7

[10] A. Doshi and M. M. Trivedi. Head and eye gaze dynamics during visual attention shifts in complex environments. *Journal of Vision*, 12(2):9, 2012. 6

[11] A. T. Duchowski, N. Cournia, and H. Murphy. Gaze-contingent displays: a review. *Cyberpsychol Behav.*, 7(6):621–34, 2004. 2

[12] M. Feng, A. Borji, and H. Lu. Fixation prediction with a combined model of bottom-up saliency and vanishing point. In *Proc. IEEE WACV*, pages 1–7. IEEE, 2016. 1

[13] E. G. Freedman. Coordination of the eyes and head during visual orienting. *Experimental brain research*, 190(4):369–387, 2008. 2, 6

[14] A. Gibaldi, M. Vanegas, P. J. Bex, and G. Maiello. Evaluation of the tobii eyex eye tracking controller and matlab toolkit for research. *Behavior Research Methods*, pages 1–24, 2016. 3

[15] S. Goferman, L. Zelnik-Manor, and A. Tal. Context-aware saliency detection. *IEEE Trans. PAMI*, 34(10):1915–1926, 2012. 2

[16] F. Guo, J. Shen, and X. Li. Learning to detect stereo saliency. In *Proc. IEEE ICME*, pages 1–6. IEEE, 2014. 2

[17] L. Itti, C. Koch, E. Niebur, et al. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. PAMI*, 20(11):1254–1259, 1998. 2

[18] Y. Jia and M. Han. Category-independent object-level saliency detection. In *Proc. IEEE CVPR*, pages 1761–1768, 2013. 2

[19] T. Judd, K. Ehinger, F. Durand, and A. Torralba. Learning to predict where humans look. In *Proc. IEEE ICCV*, 2009. 2, 3, 4

[20] W. Kienzle, F. A. Wichmann, M. O. Franz, and B. Schölkopf. A nonparametric approach to bottom-up visual saliency. In *Proc. NIPS*, pages 689–696, 2006. 2

[21] C. Koch and S. Ullman. Shifts in selective visual attention: towards the underlying neural circuitry. In *Matters of intelligence*, pages 115–141. Springer, 1987. 1

[22] T. Kollenberg, A. Neumann, D. Schneider, T.-K. Tews, T. Hermann, H. Ritter, A. Dierker, and H. Koesling. Visual search in the (un) real world: how head-mounted displays affect eye movements, head movements and target detection. In *Proc. ACM ETRA*, pages 121–124. ACM, 2010. 2, 6

[23] V. Laurutis and D. Robinson. The vestibulo-ocular reflex during human saccadic eye movements. *The Journal of Physiology*, 373(1):209–233, 1986. 2, 6

[24] O. Le Meur and T. Baccino. Methods for comparing scanpaths and saliency maps: strengths and weaknesses. *Behavior research methods*, 45(1):251–266, 2013. 3

[25] O. Le Meur, T. Baccino, and A. Roumy. Prediction of the inter-observer visual congruency (iovc) and application to image ranking. In *Proc. ACM Int. Conf. on Multimedia*, pages 373–382. ACM, 2011. 3

[26] G. Leifman, D. Rudoy, T. Swedish, E. Bayro-Corrochano, and R. Raskar. Learning gaze transitions from depth to improve video saliency estimation. *arXiv preprint arXiv:1603.03669*, 2016. 2

[27] G. Li and Y. Yu. Visual saliency based on multiscale deep features. In *Proc. IEEE CVPR*, pages 5455–5463, 2015. 2

[28] R. Liu, J. Cao, Z. Lin, and S. Shan. Adaptive partial differential equation learning for visual saliency detection. In *Proc. IEEE CVPR*, pages 3866–3873, 2014. 2

[29] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, X. Tang, and H.-Y. Shum. Learning to detect a salient object. *IEEE Trans. PAMI*, 33(2):353–367, 2011. 2

[30] R. Nakashima, Y. Fang, Y. Hatori, A. Hiratani, K. Matsumiya, I. Kuriki, and S. Shioiri. Saliency-based gaze prediction based on head direction. *Vision research*, 117:59–66, 2015. 2

[31] A. Nuthmann and J. M. Henderson. Object-based attentional selection in scene viewing. *Journal of Vision*, 10(8):20, 2010. 4

[32] J. Pan, E. Sayrol, X. Giro-i Nieto, K. McGuinness, and N. E. O'Connor. Shallow and deep convolutional networks for saliency prediction. In *Proc. IEEE CVPR*, June 2016. 2, 7

[33] K. Ruhland, C. E. Peters, S. Andrist, J. B. Badler, N. I. Badler, M. Gleicher, B. Mutlu, and R. McDonnell. A review of eye gaze in virtual agents, social robotics and hci: Behaviour generation, user interaction and perception. In *Computer Graphics Forum*, volume 34, pages 299–326, 2015. 2

[34] D. D. Salvucci and J. H. Goldberg. Identifying fixations and saccades in eye-tracking protocols. In *Proc. ACM ETRA*, pages 71–78. ACM, 2000. 3, 6

[35] V. Tanriverdi and R. J. K. Jacob. Interacting with eye movements in virtual environments. In *SIGCHI*, pages 265–272, 2000. 2

[36] A. Torralba, A. Oliva, M. S. Castelhano, and J. M. Henderson. Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychological review*, 113(4):766, 2006. 2

[37] E. Vig, M. Dorr, and D. Cox. Large-scale optimization of hierarchical features for saliency prediction in natural images. In *Proc. IEEE CVPR*, pages 2798–2805, 2014. 2

[38] L. Wang, H. Lu, X. Ruan, and M.-H. Yang. Deep networks for saliency detection via local estimation and global search. In *Proc. IEEE CVPR*, pages 3183–3192, 2015. 2

[39] M. Yu, H. Lakshman, and B. Girod. A framework to evaluate omnidirectional video coding schemes. In *Proc. ISMAR*, 2015. 8

[40] Q. Zhao and C. Koch. Learning saliency-based visual attention: A review. *Signal Processing*, 93(6):1401–1407, 2013. 2

[41] R. Zhao, W. Ouyang, H. Li, and X. Wang. Saliency detection by multi-context deep learning. In *Proc. IEEE CVPR*, pages 1265–1274, 2015. 2